Intervention Mechanism Design for Networks With Selfish Users

Jaeok Park and Mihaela van der Schaar

Abstract

We consider a multi-user network where a network manager and selfish users interact. The network manager monitors the behavior of users and intervenes in the interaction among users if necessary, while users make decisions independently to optimize their individual objectives. In this paper, we develop a framework of intervention mechanism design, which is aimed to optimize the objective of the manager, or the network performance, taking the incentives of selfish users into account. Our framework is general enough to cover a wide range of application scenarios, and it has advantages over existing approaches such as Stackelberg strategies and pricing. To design an intervention mechanism and to predict the resulting operating point, we formulate a new class of games called intervention games and a new solution concept called intervention equilibrium. We provide analytic results about intervention equilibrium and optimal intervention mechanisms in the case of a benevolent manager with perfect monitoring. We illustrate these results with a random access model. Our illustrative example suggests that intervention requires less knowledge about users than pricing.

Index Terms

Game theory, incentives, intervention, mechanism design, multi-user networks, network management

I. Introduction

A. Motivation and Our Approach

In noncooperative multi-user networks, multiple users interact with each other making decisions independently to optimize their individual objectives. Operation by selfish users often results in a suboptimal network performance, and this calls for an incentive mechanism to guide selfish users to behave according

The authors are with Electrical Engineering Department, University of California, Los Angeles (UCLA), 420 Westwood Plaza, Los Angeles, CA 90095-1594, USA. e-mail: {jaeok, mihaela}@ee.ucla.edu.

to the system objective. In this paper, we propose a class of incentive mechanisms, called *intervention mechanisms*, that can be used when there is a network manager that can monitor the behavior of users and intervene in the interaction among users if necessary. Since the manager can react to the behavior of users through an intervention mechanism, an intervention mechanism has the potential to shape the incentives of users.

To design an intervention mechanism and to predict the resulting operating point, we formulate a new class of games called *intervention games* and a new solution concept called *intervention equilibrium*. In an intervention game, the manager first chooses an intervention mechanism, and then users choose their actions knowing the intervention mechanism chosen by the manager. After observing a signal about the actions of users, the manager chooses its action according to the intervention mechanism. An intervention mechanism specifies an action that the manager takes following each possible signal realization, and thus can be considered as a protocol that prescribes a rule according to which the manager rewards and punishes users depending on their observed behavior. The manager chooses an intervention mechanism to optimize its objective, anticipating the rational behavior of users given the intervention mechanism it chooses. Intervention equilibrium predicts the outcome of an intervention game in terms of an intervention mechanism designed by the manager and an operating point chosen by users.

B. Related Work

An approach similar to ours can be found in [1]. The authors of [1] consider a congestion control problem and analyze performance under different scheduling mechanisms, which can be considered as intervention mechanisms. With a scheduling mechanism that assigns a higher priority to flows with a smaller rate, the input rates of users can be restrained voluntarily. In other words, users do not send their traffic excessively in their self-interest because doing so will result in a higher probability that their packets are dropped. Another closely related work is our previous work [2]. In [2], we consider a random access network where the manager can intervene by jamming the packets of users and show that intervention mechanisms can successfully control the transmission probabilities of selfish users. However, these works consider only specific communication networks. In this paper, we aim to develop a general framework of intervention mechanisms that can be applied to various application scenarios including not only communication networks but also other scenarios such as file sharing on peer-to-peer networks, multi-task processing on cloud computers, and load balancing on multi-core processors.

In an intervention game, the manager makes a decision before users do, and thus the manager can be considered as a leader and users as followers. Such a hierarchical structure of multi-user interaction has

been modeled as a Stackelberg game in the context of congestion control [3], medium access control [4], and power control [5]. The main difference between our intervention approach and the traditional Stackelberg approach can be found in the ability of the manager. In an intervention game, the manager can monitor the actions of users and can commit to an intervention mechanism. Hence, a strategy of the manager is its complete contingent plan for actions to be taken given a signal realization. On the other hand, in the Stackelberg games in [3]–[5], the leader does not have monitoring and commitment capabilities, and thus simply chooses an action before followers choose their actions. Our formulation includes the traditional Stackelberg formulation as a special case, because an intervention mechanism reduces to an action if the manager chooses the same action regardless of signal realizations or if it observes no signal. Therefore, intervention games yield higher design flexibility for the manager than traditional Stackelberg games do, enabling the greater potential to shape the incentives of users.

Pricing mechanisms that charge users for their usage have received a significant amount of attention in the literature [6]. Pricing mechanisms have a solid theoretical foundation in economics as well as high design flexibility. They can be considered as a special class of intervention mechanisms, where the intervention of the manager corresponds to charging and collecting payments from users. The existing literature has studied pricing mechanisms using Stackelberg games in the context of congestion control [7] and cognitive radio networks [8].

C. Focus and Advantages of Our Approach

Our discussion so far suggests a classification of intervention into direct and indirect intervention. Intervention is direct (resp. indirect) if the manager interacts with users in the inside (resp. outside) of the network. In other words, the manager with direct intervention intervenes in the network usage of users, while the manager with indirect intervention influences the utilities of users through an outside instrument, for example, monetary payments in the case of pricing. Direct intervention can be further classified into adaptive and constant intervention, depending on whether the manager can react to the actions of users. Traditional Stackelberg strategies where the manager takes an action before users do belong to constant intervention. Table I classifies the existing intervention approaches in the literature according to the above criteria.

Although our framework covers all the three classes of intervention in Table I, our main focus is on adaptive direct intervention, which has not received much attention from researchers compared to the other two classes. Adaptive direct intervention is useful compared to constant direct intervention especially when the manager does not value its usage of the network. In this scenario, the manager

TABLE I

CLASSIFICATION OF THE EXISTING INTERVENTION APPROACHES IN THE LITERATURE.

Direct intervention		Indirect intervention (e.g., pricing)
Adaptive	Constant	
[1], [2]	[3], [4], [5]	[7], [8]

desires to achieve a certain operating point with the minimum level of intervention. With constant direct intervention the manager needs to consume network resources in order to impact the behavior of users, whereas with adaptive direct intervention the manager can use intervention only as a threat to punish users in case of deviation. Direct intervention has an advantage over indirect intervention in terms of implementation. Since direct intervention affects the network usage of users directly, users cannot evade intervention as long as they use the network. On the contrary, enforcing the outside instrument of indirect intervention can be costly especially when the network is freely accessible. The difference between direct and indirect intervention yields an informational advantage of direct intervention over indirect intervention. For instance, direct intervention in [1] and [2] affects the data rates of users. Hence, the effectiveness of a direct intervention mechanism is independent of the way users value their data rate. On the contrary, designing a pricing mechanism to achieve a certain operating point in general requires the manager to know the utility functions of users. This point is illustrated with an example in Section V.

D. Comparison with Other Frameworks

We first compare the framework of intervention mechanism design with that of network utility maximization (NUM) [9]. Since our framework allows a general form of the objective of the manager, the objective can be set as the sum of the utilities of users, as in NUM. The main difference between intervention and NUM is that intervention takes into account the incentives of selfish users whereas NUM assumes obedient users. The NUM framework aims to design a distributed algorithm following which obedient users can reach a system-wide optimal operating point, using prices as congestion signals. On the contrary, the intervention framework aims to design an incentive mechanism that induces selfish users to achieve an incentive-constrained optimal operating point. Hence, the intervention framework is more relevant in a network with selfish users.

Next, we compare intervention games and repeated games [10], which share a common goal of sustaining cooperation among selfish individuals. With a repeated game strategy such as tit-for-tat, users monitor the actions of other users and choose their actions depending on their past observations. Hence,

the burden of monitoring and executing reward and punishment is distributed to users in a repeated game, while it is imposed solely on the manager in an intervention game. If there is no reliable manager in the network, a mechanism based on a repeated game strategy is a possible alternative to an intervention mechanism. However, in order to implement a repeated game strategy, users in the network must interact frequently and maintain histories of their observations. Moreover, sustaining an optimal operating point with a repeated game strategy often requires users to be sufficiently patient. On the contrary, implementing intervention mechanisms effectively requires neither repeated interaction nor patience of users.

Lastly, we compare intervention mechanism design with standard mechanism design [11], [12]. A standard mechanism design problem considers a scenario where the manager can control the system configuration but has incomplete information about the types of users. Thus, standard mechanism design is concerned about incentives regarding types, whereas intervention mechanism design in this paper cares about incentives regarding actions. In economics terminology, standard mechanism design focuses on *adverse selection* while our intervention mechanism design focuses on *moral hazard* [11]. We can modify or extend our framework to study the capabilities and limitations of intervention mechanisms to overcome incentive problems in the presence of only adverse selection or both adverse selection and moral hazard. We leave these topics for future research.

E. Organization of the Paper

The remainder of this paper is organized as follows. In Section II, we formulate the framework of intervention mechanism design, defining intervention games and intervention equilibrium. In Section III, we classify the types of intervention and the manager to point out the broad scope of our framework. In Section IV, we consider a benevolent manager with perfect monitoring and provide analytic results about intervention equilibrium. In Section V, we apply our framework to random access networks in order to illustrate the results. We conclude in Section VI.

II. Framework

We consider a multi-user network where users interact with each other taking actions independently. There is a network manager that can intervene in the interaction of users. There are N users in the network, and the set of users is denoted by $\mathcal{N} = \{1, \ldots, N\}$. For convenience, we call the manager user 0. To distinguish the users from the manager, we sometimes refer to a user as a regular user. The regular users are indexed from user 1 to user N. The set of all the users including the manager is denoted by $\mathcal{N}_0 = \mathcal{N} \cup \{0\}$. Each user $i \in \mathcal{N}_0$ chooses an action a_i from the set of actions available to it, which

is denoted by A_i. An action profile of the regular users is written as $a = (a_1, \ldots, a_N)$, while the set of action profiles of the regular users is written as $A = \prod_{i \in \mathcal{N}} A_i$. An action profile of the regular users other than regular user i is written as $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_N)$ so that a can be expressed as $a = (a_i, a_{-i})$.

The actions of the manager and the users jointly determine utilities they receive. The utility function of user $i \in \mathcal{N}_0$ is denoted by $u_i : A_0 \times A \to \mathbb{R}$. That is, $u_i(a_0, a)$ represents the utility that user i receives when the manager chooses action a_0 and the users choose an action profile a. The manager is different from the regular users in that the manager is able to monitor the actions of the regular users before it takes its action. The monitoring ability of the manager is formally represented by a monitoring technology (S, ρ). S is the set of all possible signals that the manager can obtain, while ρ is a mapping from A to $\triangle(S)$, where $\triangle(S)$ is the set of probability distributions over S. That is, $\rho(a)$ represents the probability distribution of signals when the regular users choose a.

Since the manager takes its action after observing a signal, a strategy for the manager is its complete contingent plan for actions to be taken following all possible signal realizations while a strategy for a regular user is simply its choice of an action.[1] Thus, a strategy for the manager can be represented by a function $f : S \to A_0$, which we call an *intervention mechanism*. That is, an intervention mechanism specifies an action taken by the manager following each signal it can observe. Let $\mathcal{F}$ be the set of all intervention mechanisms, i.e., the set of all functions from S to A_0. We assume that the manager is able to commit to an intervention mechanism and that the intervention mechanism chosen by the manager is known to the regular users when they choose their actions.[2]

To sum up, an *intervention game* is summarized by the data

$$\mathcal{G} = \langle \mathcal{N}_0, (A_i)_{i \in \mathcal{N}_0}, (u_i)_{i \in \mathcal{N}_0}, (S, \rho) \rangle ,$$

and the timing of an intervention game can be described as follows.

1) The manager chooses an intervention mechanism $f \in \mathcal{F}$.
2) The regular users choose their actions $a \in A$ independently and simultaneously, knowing the intervention mechanism f chosen by the manager.
3) The manager observes a signal $s \in S$, which is realized following a probability distribution $\rho(a) \in \triangle(S)$.

[1]For expositional simplicity, we restrict attention to pure strategies, although our formulation extends easily to the case of randomized strategies.

[2]The information about the intervention mechanism can be obtained by an agreed protocol specification or by learning.

4) The manager chooses its action $a_0 \in A_0$ according to f, i.e., $a_0 = f(s)$.
5) The users receive their utilities based on the chosen actions $(a_0, a) \in A_0 \times A$.

Consider an intervention mechanism $f \in \mathcal{F}$, and define a function $v_i^f : A \to \mathbb{R}$ by

$$v_i^f(a) = E_{\rho(a)}[u_i(f(s), a)], \tag{1}$$

for all $i \in \mathcal{N}$, where $E_{\rho(a)}[\cdot]$ denotes expectation with respect to the random variable s following the distribution $\rho(a)$. Then $v_i^f(a)$ represents the expected utility that regular user i receives when the manager uses intervention mechanism f and the regular users choose action profile a. An intervention mechanism f induces a simultaneous game played by the regular users, whose normal form representation is given by

$$\mathcal{G}_f = \left\langle \mathcal{N}, (A_i)_{i \in \mathcal{N}}, (v_i^f)_{i \in \mathcal{N}} \right\rangle.$$

We can predict actions chosen by the selfish regular users given an intervention mechanism f by applying the solution concept of Nash equilibrium to the induced game $\mathcal{G}_f$.

Definition 1: An intervention mechanism $f \in \mathcal{F}$ *supports* the action profile $a^* \in A$ of the regular users if a^* is a Nash equilibrium of the game $\mathcal{G}_f$, i.e.,

$$v_i^f(a_i^*, a_{-i}^*) \geq v_i^f(a_i, a_{-i}^*) \quad \text{for all } a_i \in A_i, \text{ for all } i \in \mathcal{N}.$$

Also, f *strongly supports* a^* if a^* is a unique Nash equilibrium of $\mathcal{G}_f$.

If an action profile a^* is supported by an intervention mechanism f, the regular users cannot gain by a unilateral deviation from a^* as long as the manager uses intervention mechanism f. When choosing an intervention mechanism, the manager can expect that the regular users will choose an action profile that is supported by the intervention mechanism it chooses. However, when there are multiple action profiles supported by the intervention mechanism, the manager may be uncertain about the action profile chosen by the regular users. This uncertainty disappears when the intervention mechanism strongly supports an action profile. In the formulation of our solution concept, we assume that the regular users will always select the best Nash equilibrium for the manager if the induced game has multiple Nash equilibria.

Definition 2: $(f^*, a^*) \in \mathcal{F} \times A$ is an (strong) equilibrium of intervention game $\mathcal{G}$, or an *(strong) intervention equilibrium,* if f^* (strongly) supports a^* and

$$E_{\rho(a^*)}[u_0(f^*(s), a^*)] \geq E_{\rho(a)}[u_0(f(s), a)] \quad \text{for all } (f, a) \in \mathcal{F} \times A \text{ such that } f \text{ supports } a.$$

$f^* \in \mathcal{F}$ is an *(strongly) optimal intervention mechanism* if there exists an action profile $a^* \in A$ such that (f^*, a^*) is an (strong) intervention equilibrium.

Intervention equilibrium is a solution concept for intervention games, based on a backward induction argument, assuming that the manager can commit to an intervention mechanism and predict the rational reaction of the regular users to its choice of an intervention mechanism. An intervention equilibrium can be considered as a Stackelberg equilibrium (or a subgame perfect equilibrium) applied to an intervention game $\mathcal{G}$, since the induced game $\mathcal{G}_f$ is a subgame of $\mathcal{G}$. Since we assume that the manager can induce the regular users to choose the best Nash equilibrium for it in the case of multiple Nash equilibria, the problem of designing an optimal intervention mechanism can be expressed as $\max_{f \in \mathcal{F}} \max_{a \in \mathcal{E}(f)} E_{\rho(a)}[u_0(f(s), a)]$, where $\mathcal{E}(f)$ denotes the set of action profiles supported by f.

III. Classification of Intervention and the Manager

In this section, we classify the types of intervention depending on the form of the utility functions of the regular users (i.e., the way the manager interacts with the regular users) and the types of the manager depending on the form of the utility function the manager (i.e., the objective of the manager). The purpose of classification is to suggest that a wide range of application scenarios can be modeled in the framework of intervention.

Definition 3: Let $A_i \subset \mathbb{R}^K$ for some K, for all $i \in \mathcal{N}_0$. Intervention is *additively symmetric* if, for each $i \in \mathcal{N}$, there exist functions $g_{ij}^k : \mathbb{R} \to \mathbb{R}$ and $h_i^k : \mathbb{R}^2 \to \mathbb{R}$, for $k = 1, \ldots, K$ and for $j \in \mathcal{N}_0 \setminus \{i\}$, such that

$$u_i(a_0, a) = \sum_{k=1}^{K} h_i^k \left(a_i^k, \sum_{j \in \mathcal{N}_0 \setminus \{i\}} g_{ij}^k(a_j^k) \right),$$

where a_i^k denotes the k-th element of a_i. Intervention is *multiplicatively symmetric* if, for each $i \in \mathcal{N}$, there exist functions $g_{ij}^k : \mathbb{R} \to \mathbb{R}$ and $h_i^k : \mathbb{R}^2 \to \mathbb{R}$, for $k = 1, \ldots, K$ and for $j \in \mathcal{N}_0 \setminus \{i\}$, such that

$$u_i(a_0, a) = \sum_{k=1}^{K} h_i^k \left(a_i^k, \prod_{j \in \mathcal{N}_0 \setminus \{i\}} g_{ij}^k(a_j^k) \right).$$

Intervention is *symmetric* if it is either additively or multiplicatively symmetric.

The manager with symmetric intervention uses the network as the regular users do, in the same position as the regular users. The effect of the actions of other uses including the manager on a regular user is represented by an additive term in the case of additive symmetric intervention and a multiplicative term in the case of multiplicatively symmetric intervention. Network models in which the manager can exert additive symmetric intervention include routing [3] and frequency-selective Gaussian interference channels [5]. Multiplicative symmetric intervention can be used in the random access model of [2].

We can interpret the objective of the manager as the system objective, which can vary depending on the types of the manager. Below we classify the manager with symmetric intervention depending on the form of its utility function.

Definition 4: Consider symmetric intervention. The manager is *benevolent* if

$$u_0(a_0, a) = \sum_{i \in \mathcal{N}} w_i u_i(a_0, a)$$

for some $(w_1, \ldots, w_N) \in \mathbb{R}_+^N$. The manager is *self-interested* if there exist functions $g_{0j}^k : \mathbb{R} \to \mathbb{R}$ and $h_0^k : \mathbb{R}^2 \to \mathbb{R}$, for $k = 1, \ldots, K$ and for $j \in \mathcal{N}$, such that

$$u_0(a_0, a) = \sum_{k=1}^{K} h_0^k \left(a_i^k, \sum_{j \in \mathcal{N}} g_{0j}^k(a_j^k) \right) \tag{2}$$

in the case of additively symmetric intervention, and

$$u_0(a_0, a) = \sum_{k=1}^{K} h_0^k \left(a_i^k, \prod_{j \in \mathcal{N}} g_{0j}^k(a_j^k) \right) \tag{3}$$

in the case of multiplicatively symmetric intervention. The manager is *total welfare maximizing* if

$$u_0(a_0, a) = w_0 \tilde{u}_0(a_0, a) + \sum_{i \in \mathcal{N}} w_i u_i(a_0, a)$$

for some $(w_0, w_1, \ldots, w_N) \in \mathbb{R}_+^{N+1}$, where $\tilde{u}_0(a_0, a)$ denotes the right-hand side of (2) if intervention is additively symmetric and that of (3) if intervention is multiplicatively symmetric.

A benevolent manager maximizes the welfare of the regular users, attaching a welfare weight w_i to regular user i. The utility function of a self-interested manager has the same structure as those of the regular users. Recall that the manager with symmetric intervention uses the network as the regular users do. A benevolent manager does not derive any utility from its usage of the network, whereas the utility of a self-interested manager is derived solely from its usage [5]. A total welfare maximizing manager derives utility from its usage as well as that of the other users [3].

It is also possible that the manager plays a special role in the network, participating in interaction in a different way from the regular users. For example, the manager plays the role of a scheduler in [1] and a billing authority in pricing [7], [8]. We call this type of intervention asymmetric intervention to contrast it with symmetric intervention.[3] We present two representative examples of asymmetric intervention, which arise naturally in games with transferable utility [13].

[3]Indirect intervention is necessarily asymmetric, while direct intervention can be symmetric or asymmetric.

Definition 5: Intervention is *asymmetric* if it is not symmetric. Let $A_0 = \mathbb{R}_+^N$. Intervention is *additively asymmetric* if, for each $i \in \mathcal{N}$, there exists function $g_i : A \to \mathbb{R}$ such that

$$u_i(a_0, a) = g_i(a) - a_0^i,$$

where a_0^i is the i-th element of a_0. Let $A_0 = [0,1]^N$. Intervention is *multiplicatively asymmetric* if, for each $i \in \mathcal{N}$, there exists function $g_i : A \to \mathbb{R}_+$ such that

$$u_i(a_0, a) = (1 - a_0^i) g_i(a).$$

With additively asymmetric intervention, the manager takes away the amount a_0^i from the benefit $g_i(a)$ that regular user i receives from its usage of the network. With multiplicatively asymmetric intervention, the manager deducts the a_0^i fraction of the benefit that regular user i receives. Hence, the two types of asymmetric intervention in Definition 5 can be compared to two different forms of taxation, where additively asymmetric intervention corresponds to taxation with lump-sum tax rates and multiplicatively asymmetric intervention to taxation with proportional tax rates.

Definition 6: Consider additively or multiplicatively asymmetric intervention. Let $c_0 : A \to \mathbb{R}_+$ be a function that represents the operating cost of the network. The manager is *benevolent* if

$$u_0(a_0, a) = \sum_{i \in \mathcal{N}} u_i(a_0, a).$$

The manager is *self-interested* if

$$u_0(a_0, a) = \sum_{i \in \mathcal{N}} a_0^i - c_0(a)$$

in the case of additively asymmetric intervention, and

$$u_0(a_0, a) = \sum_{i \in \mathcal{N}} a_0^i g_i(a) - c_0(a)$$

in the case of multiplicatively asymmetric intervention. The manager is *total welfare maximizing* if

$$u_0(a_0, a) = \sum_{i \in \mathcal{N}} g_i(a) - c_0(a).$$

In Definition 6, we use welfare weights $w_i = 1$ for all $i \in \mathcal{N}_0$ because of the assumption of transferable utility. For ease of interpretation, let us regard the transfer from the regular users to the manager as payments. A benevolent manager does not value payments (or burns payments) it receives from the regular users, and thus payments create a welfare loss. On the contrary, a self-interested manager values payments from the regular users, and it maximizes its profit measured by the total payment minus the operating cost. If there is no operating cost, i.e., $c_0 \equiv 0$, then the objective of a self-interested manager is revenue

maximization. A total welfare maximizing manager maximizes the net gain from operating the network, which is the total benefit of the regular users minus the operating cost. We can also consider an *individually rational benevolent* manager that has utility function $u_0(a_0, a) = \sum_{i \in \mathcal{N}} u_i(a_0, a)$ and faces the individual rationality constraint that requires the total payment it receives to be no less than the operating cost (i.e., $\sum_{i \in \mathcal{N}} a_0^i \geq c_0(a)$ in the case of additively asymmetric intervention and $\sum_{i \in \mathcal{N}} a_0^i g_i(a) \geq c_0(a)$ in the case of multiplicatively asymmetric intervention). In case that a regular user has an outside option, the individual rationality constraint for the regular user can be taken care of by including in its action space an action that corresponds to choosing its outside option.

IV. Benevolent Manager with Perfect Monitoring

In this section, we analyze a class of intervention games that satisfy the following maintained assumptions, while leaving the analysis of other classes for future work.

Assumption 1: (i) (Benevolent manager) The utility function of the manager is given by $u_0(a_0, a) = \sum_{i \in \mathcal{N}} u_i(a_0, a)$ for all $(a_0, a) \in A_0 \times A$.
(ii) (Existence of minimal and maximal intervention actions) There exist the minimal and maximal elements of A_0, denoted $\underline{a}_0$ and $\overline{a}_0$, respectively, in the sense that for all $i \in \mathcal{N}$, $\underline{a}_0$ and $\overline{a}_0$ satisfy

$$u_i(\underline{a}_0, a) \geq u_i(a_0, a) \geq u_i(\overline{a}_0, a) \quad \text{for all } a_0 \in A_0, \text{ for all } a \in A.$$

(iii) (Perfect monitoring) The monitoring technology of the manager is perfect in the sense that the manager can observe the actions of the regular users without errors. Formally, a monitoring technology (S, ρ) is perfect if $S = A$ and only signal a can arise in the distribution $\rho(a)$ for all $a \in A$.

In Assumption 1(i), we set welfare weights as $w_i = 1$ for all $i \in \mathcal{N}$ in order to cover both cases of transferable and nontransferable utility, although our results extend easily to general welfare weights. In Assumption 1(ii), $\underline{a}_0$ and $\overline{a}_0$ can be interpreted as the minimal and maximal intervention actions of the manager, respectively. For given $a \in A$, each regular user receives the highest (resp. lowest) utility when the manager takes the minimal (resp. maximal) intervention action. In other words, the utilities of the regular users are aligned with respect to the action of the manager so that the manager can reward or punish all the regular users at the same time. Combining Assumption 1(i) and (ii), we obtain

$$u_0(\underline{a}_0, a) \geq u_0(a_0, a) \geq u_0(\overline{a}_0, a) \quad \text{for all } a_0 \in A_0, \text{ for all } a \in A. \tag{4}$$

Thus, for given $a \in A$ the benevolent manager prefers to use the minimal intervention action. With perfect monitoring, v_i^f defined in (1) reduces to $v_i^f(a) = u_i(f(a), a)$ for all $a \in A$.

We first characterize the set of action profiles of the regular users that can be supported by an intervention mechanism. Define $\mathcal{E} = \cup_{f \in \mathcal{F}} \mathcal{E}(f) = \{a \in A : \exists f \in \mathcal{F}$ such that f supports $a\}$.

Proposition 1: $a^* \in \mathcal{E}$ if and only if $u_i(\underline{a}_0, a^*) \geq u_i(\overline{a}_0, a_i, a^*_{-i})$ for all $a_i \in A_i$, for all $i \in \mathcal{N}$.

Proof: Suppose that $u_i(\underline{a}_0, a^*) \geq u_i(\overline{a}_0, a_i, a^*_{-i})$ for all $a_i \in A_i$, for all $i \in \mathcal{N}$. Define an intervention mechanism $f_{\tilde{a}}$, for each $\tilde{a} \in A$, by

$$f_{\tilde{a}}(a) = \begin{cases} \underline{a}_0 & \text{if } a = \tilde{a}, \\ \overline{a}_0 & \text{otherwise.} \end{cases} \tag{5}$$

Then f_{a^*} supports a^*, and thus $a^* \in \mathcal{E}$.

Suppose that $a^* \in \mathcal{E}$. Then there exists an intervention mechanism f such that $u_i(f(a^*), a^*) \geq u_i(f(a_i, a^*_{-i}), a_i, a^*_{-i})$ for all $a_i \in A_i$, for all $i \in \mathcal{N}$. Then we obtain $u_i(\underline{a}_0, a^*) \geq u_i(f(a^*), a^*) \geq u_i(f(a_i, a^*_{-i}), a_i, a^*_{-i}) \geq u_i(\overline{a}_0, a_i, a^*_{-i})$ for all $a_i \in A_i$, for all $i \in \mathcal{N}$, where the first and the third inequalities follow from Assumption 1(ii). ■

The basic idea underlying Proposition 1 is that for given $a^* \in A$, f_{a^*} is the most effective intervention mechanism to support a^*. Thus, in order to find out whether a^* is supported by some intervention function, it suffices to check whether a^* is supported by f_{a^*}. We call $f_{\tilde{a}}$, defined in (5), the *maximum punishment* intervention mechanism with target action profile $\tilde{a}$, because the manager using $f_{\tilde{a}}$ takes the maximum intervention action whenever the users do not follow the action profile $\tilde{a}$. Let $\mathcal{F}^p$ be the set of all maximum punishment intervention mechanisms, i.e., $\mathcal{F}^p = \{f_{\tilde{a}} \in \mathcal{F} : \tilde{a} \in A\}$. Also, define $\mathcal{E}^p = \cup_{f \in \mathcal{F}^p} \mathcal{E}(f) = \{a \in A : \exists f \in \mathcal{F}^p$ such that f supports $a\}$. The second part of the proof of Proposition 1 shows that if a^* is supported by some intervention mechanism f, then it is also supported by the maximum punishment intervention mechanism with target action profile a^*, f_{a^*}. This observation leads us to the following corollary.

Corollary 1: (i) $\mathcal{E} = \mathcal{E}^p$.

(ii) If (f^*, a^*) is an intervention equilibrium, then (f_{a^*}, a^*) is also an intervention equilibrium.

Proof: (i) $\mathcal{E} \supset \mathcal{E}^p$ follows from $\mathcal{F} \supset \mathcal{F}^p$, while $\mathcal{E} \subset \mathcal{E}^p$ follows from Proposition 1.

(ii) Suppose that (f^*, a^*) is an intervention equilibrium. Then by Definition 2, f^* supports a^*, and $u_0(f^*(a^*), a^*) \geq u_0(f(a), a)$ for all $(f, a) \in \mathcal{F} \times A$ such that f supports a. Since $a^* \in \mathcal{E}$, f_{a^*} supports a^* by Proposition 1. Hence, $u_0(f^*(a^*), a^*) \geq u_0(f_{a^*}(a^*), a^*)$. On the other hand, since $f_{a^*}(a^*) = \underline{a}_0$, we have $u_0(f^*(a^*), a^*) \leq u_0(f_{a^*}(a^*), a^*)$ by (4). Therefore, $u_0(f^*(a^*), a^*) = u_0(f_{a^*}(a^*), a^*)$, and thus $u_0(f_{a^*}(a^*), a^*) \geq u_0(f(a), a)$ for all $(f, a) \in \mathcal{F} \times A$ such that f supports a. This proves that (f_{a^*}, a^*) is an intervention equilibrium. ■

Corollary 1 shows that there is no loss of generality in two senses when we restrict attention to maximum punishment intervention mechanisms. First, the set of action profiles that can be supported by an intervention mechanism remains the same when we consider only maximum punishment intervention mechanisms. Second, if there exists an optimal intervention mechanism, we can find an optimal intervention mechanism among maximum punishment intervention mechanisms. The following proposition provides a necessary and sufficient condition under which a maximum punishment intervention mechanism together with its target action profile constitutes an intervention equilibrium.

Proposition 2: (f_{a^*}, a^*) is an intervention equilibrium if and only if $a^* \in \mathcal{E}$ and $u_0(\underline{a}_0, a^*) \geq u_0(\underline{a}_0, a)$ for all $a \in \mathcal{E}$.

Proof: Suppose that (f_{a^*}, a^*) is an intervention equilibrium. Then f_{a^*} supports a^*, and thus $a^* \in \mathcal{E}$. Also, $u_0(f_{a^*}(a^*), a^*) \geq u_0(f(a), a)$ for all $(f, a) \in \mathcal{F} \times A$ such that f supports a. Choose any $a \in \mathcal{E}$. Then by Proposition 1, f_a supports a, and thus $u_0(\underline{a}_0, a^*) = u_0(f_{a^*}(a^*), a^*) \geq u_0(f_a(a), a) = u_0(\underline{a}_0, a)$.

Suppose that $a^* \in \mathcal{E}$ and $u_0(\underline{a}_0, a^*) \geq u_0(\underline{a}_0, a)$ for all $a \in \mathcal{E}$. To prove that (f_{a^*}, a^*) is an intervention equilibrium, we need to show (i) f_{a^*} supports a^*, and (ii) $u_0(f_{a^*}(a^*), a^*) \geq u_0(f(a), a)$ for all $(f, a) \in \mathcal{F} \times A$ such that f supports a. Since $a^* \in \mathcal{E}$, (i) follows from Proposition 1. To prove (ii), choose any $(f, a) \in \mathcal{F} \times A$ such that f supports a. Then $u_0(f_{a^*}(a^*), a^*) = u_0(\underline{a}_0, a^*) \geq u_0(\underline{a}_0, a) \geq u_0(f(a), a)$, where the first inequality follows from $a \in \mathcal{E}$. ■

Proposition 2 implies that if we obtain an action profile a^* such that $a^* \in \arg\max_{a \in \mathcal{E}} u_0(\underline{a}_0, a)$, we can use it to construct an intervention equilibrium and thus an optimal intervention mechanism. Corollary 1(ii) implies that, when we want to find out whether a given action profile can be supported by an optimal intervention mechanism, we can consider only the maximum punishment intervention mechanism having the action profile as its target action profile. However, when we are given an optimal intervention mechanism f_{a^*} in the class of maximum punishment intervention mechanisms, it is not certain whether its target action profile a^* or some other action profile constitutes an intervention equilibrium together with f_{a^*}. The following proposition provides a sufficient condition under which a given optimal intervention mechanism f_{a^*} must be paired with its target action profile a^* to form an intervention equilibrium.

Proposition 3: Suppose that $u_0(\underline{a}_0, a) > u_0(\overline{a}_0, a)$ for all $a \in \mathcal{E}$. If $f_{a^*} \in \mathcal{F}^p$ is an optimal intervention mechanism, then (f_{a^*}, a^*) is an intervention equilibrium and there exists no other $a \neq a^*$ such that (f_{a^*}, a) is an intervention equilibrium.

Proof: Suppose that $f_{a^*} \in \mathcal{F}^p$ is an optimal intervention mechanism. Then there must exist $a' \in A$ such that (f_{a^*}, a') is an intervention equilibrium, i.e., (i) f_{a^*} supports a', and (ii) $u_0(f_{a^*}(a'), a') \geq$

$u_0(f(a), a)$ for all $(f, a) \in \mathcal{F} \times A$ such that f supports a. Suppose that there exists $a' \neq a^*$ that satisfies (i) and (ii). Since f_{a^*} supports a', we have $a' \in \mathcal{E}$ and thus $f_{a'}$ supports a'. Then, by (ii), $u_0(\overline{a}_0, a') = u_0(f_{a^*}(a'), a') \geq u_0(f_{a'}(a'), a') = u_0(\underline{a}_0, a')$, which contradicts the assumption that $u_0(\underline{a}_0, a) > u_0(\overline{a}_0, a)$ for all $a \in \mathcal{E}$. Therefore, there cannot exist $a' \neq a^*$ that satisfies (i) and (ii). Since there must exist $a' \in A$ that satisfies (i) and (ii), a^* must satisfy (i) and (ii). ■

Although maximum punishment intervention mechanisms have a simple structure and are most effective in supporting a given action profile, they also have weaknesses. First, even when (f_{a^*}, a^*) is an intervention equilibrium, there may be other action profiles that are also supported by f_{a^*}. For example, in the case of multiplicatively asymmetric intervention where we have $\underline{a}_0 = (0, \ldots, 0)$ and $\overline{a}_0 = (1, \ldots, 1)$, f_{a^*} supports any action profile a' that has at least two different elements from those of a^*. At such an action profile, the regular users receive zero utility. As long as $u_0(\underline{a}_0, a^*) > 0$, a^* Pareto dominates other action profiles that are supported by f_{a^*}, and this can provide a rationale that the regular users will select a^* among multiple Nash equilibria of the game induced by f_{a^*}. However, in principle, the regular users may select any Nash equilibrium, and the manager cannot guarantee that the regular users will choose the intended target action profile a^* when there are other Nash equilibria. Strongly optimal intervention mechanisms have robustness in that they yield a unique Nash equilibrium and thus the issue of multiple Nash equilibria does not arise. As mentioned in Section II, an implicit assumption underlying the concept of intervention equilibrium is that the regular users will always select the best Nash equilibrium for the manager in the induced game. If the manager takes a conservative approach, it may assume that the regular users will choose the worst Nash equilibrium for it. In this case, an optimal intervention mechanism solves $\max_{f \in \mathcal{F}} \min_{a \in \mathcal{E}(f)} E_{\rho(a)}[u_0(f(s), a)]$, which has a similar spirit as maximin strategies [13] in non-zero-sum games.

Another weakness of maximum punishment intervention mechanisms is that they may incur a large efficiency loss when there are errors in the system. For example, when a regular user chooses an action different from the one it intends to take by mistake (i.e., trembling hand) or when the manager receives an incorrect signal (i.e., noisy observation), the maximal intervention action is applied even if the regular users (intend to) choose the target action profile. The case of noisy observation can be covered by modeling the monitoring technology of the manager as imperfect monitoring. In order to overcome this weakness in the case of perfect monitoring, we can consider a class of continuous intervention mechanisms (i.e., intervention mechanisms represented by a continuous function from A to A_0) if the action spaces are

continua.[4] To obtain a concrete result, we consider an intervention game where $A_i = [\underline{a}_i, \overline{a}_i] \subset \mathbb{R}$ with $\underline{a}_i < \overline{a}_i$ for all $i \in \mathcal{N}_0$ and $u_i(a_0, a)$ is strictly decreasing in a_0 on $[\underline{a}_0, \overline{a}_0]$ for all $a \in A$, for all $i \in \mathcal{N}$. We define an intervention mechanism $f_{\tilde{a},c}$, for each $\tilde{a} \in A$ and $c \in \mathbb{R}^N$, by

$$f_{\tilde{a},c}(a) = [c \cdot (a - \tilde{a})]_{\underline{a}_0}^{\overline{a}_0},$$

for all $a \in A$, where $[x]_\alpha^\beta = \min\{\max\{x, \alpha\}, \beta\}$. We call $f_{\tilde{a},c}$ the *affine* intervention mechanism with target action profile $\tilde{a}$ and intervention rate profile c. The following proposition constructs an affine intervention mechanism to support an interior target action profile in the case of differentiable utility functions.

Proposition 4: Suppose that u_i is twice continuously differentiable for all $i \in \mathcal{N}$. Let $a^* \in A$ be an action profile such that $a_i^* \in (\underline{a}_i, \overline{a}_i)$ for all $i \in \mathcal{N}$, and let

$$c_i^* = -\frac{\partial u_i(\underline{a}_0, a^*)/\partial a_i}{\partial u_i(\underline{a}_0, a^*)/\partial a_0} \tag{6}$$

for all $i \in \mathcal{N}$.[5] Suppose that

$$\frac{\partial^2 u_i}{\partial a_i^2}(\underline{a}_0, a_i, a_{-i}^*) \leq 0 \quad \text{for all } a_i \in (\underline{a}_i, \overline{a}_i) \tag{7}$$

for all $i \in \mathcal{N}$ such that $c_i^* = 0$,

$$\frac{\partial^2 u_i}{\partial a_i^2}(\underline{a}_0, a_i, a_{-i}^*) \leq 0 \quad \text{for all } a_i \in (\underline{a}_i, a_i^*), \tag{8}$$

$$\left((c_i^*)^2 \frac{\partial^2 u_i}{\partial a_0^2} + 2c_i^* \frac{\partial^2 u_i}{\partial a_i \partial a_0} + \frac{\partial^2 u_i}{\partial a_i^2} \right) \bigg|_{(a_0, a_i, a_{-i}) = (c_i^*(a_i - a_i^*) + \underline{a}_0, a_i, a_{-i}^*)} \leq 0$$
$$\text{for all } a_i \in (a_i^*, \min\{\overline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*\}), \tag{9}$$

$$\frac{\partial u_i}{\partial a_i}(\overline{a}_0, a_i, a_{-i}^*) \leq 0 \quad \text{for all } a_i \in (a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*, \overline{a}_i) \tag{10}$$

for all $i \in \mathcal{N}$ such that $c_i^* > 0$, and

$$\frac{\partial u_i}{\partial a_i}(\overline{a}_0, a_i, a_{-i}^*) \geq 0 \quad \text{for all } a_i \in (\underline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*),$$

$$\left((c_i^*)^2 \frac{\partial^2 u_i}{\partial a_0^2} + 2c_i^* \frac{\partial^2 u_i}{\partial a_i \partial a_0} + \frac{\partial^2 u_i}{\partial a_i^2} \right) \bigg|_{(a_0, a_i, a_{-i}) = (c_i^*(a_i - a_i^*) + \underline{a}_0, a_i, a_{-i}^*)} \leq 0$$
$$\text{for all } a_i \in (\max\{\overline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*\}, a_i^*), \tag{11}$$

[4]Following [10], we say that an action space is a continuum if it is a compact and convex subset of the Euclidean space $\mathbb{R}^K$ for some K.

[5]We define $\partial u_i(\underline{a}_0, a^*)/\partial a_0$ as the right partial derivative of u_i with respect to a_0 at $(\underline{a}_0, a^*)$.

$$\frac{\partial^2 u_i}{\partial a_i^2}(\underline{a}_0, a_i, a_{-i}^*) \leq 0 \quad \text{for all } a_i \in (a_i^*, \overline{a}_i)$$

for all $i \in \mathcal{N}$ such that $c_i^* < 0$.[6] Then f_{a^*,c^*} supports a^*.

Proof: See the Appendix. ■

Note that $\partial u_i(\underline{a}_0, a^*)/\partial a_0 < 0$ for all $i \in \mathcal{N}$ since u_i is strictly decreasing in a_0. Thus, c_i^*, defined in (6), has the same sign as $\partial u_i(\underline{a}_0, a^*)/\partial a_i$. With $A_0 = [\underline{a}_0, \overline{a}_0]$, the action of the manager can be interpreted as the intervention level, and the regular users receive higher utility as the intervention level is smaller. The affine intervention mechanism f_{a^*,c^*}, constructed in Proposition 4, has the properties that the manager chooses the minimal intervention level $\underline{a}_0$ when the regular users choose the target action profile a^*, i.e., $f_{a^*,c^*}(a^*) = \underline{a}_0$, and that the intervention level increases in the rate of $|c_i^*|$ as regular user i deviates to the direction in which its utility increases at $(\underline{a}_0, a^*)$. The expression of c_i^* in (6) has an intuitive explanation. Since c_i^* is proportional to $\partial u_i(\underline{a}_0, a^*)/\partial a_i$ and inversely proportional to $-\partial u_i(\underline{a}_0, a^*)/\partial a_0$, a regular user faces a higher intervention rate as its incentive to deviate from $(\underline{a}_0, a^*)$ is stronger and as a change in the intervention level has a smaller impact on its utility. The intervention level does not react to the action of regular user i when $c_i^* = 0$, because regular user i chooses a_i^* in its self-interest even when the intervention level is fixed at $\underline{a}_0$, provided that other regular users choose a_{-i}^*. Finally, we note that if (f_{a^*}, a^*) is an intervention equilibrium and $f_{a^*,c}$ supports a^* for some c, then $(f_{a^*,c}, a^*)$ is also an intervention equilibrium, since $f_{a^*}(a^*) = f_{a^*,c}(a^*) = \underline{a}_0$.

V. Application to Random Access Networks

In this section, we illustrate the results of Section IV by introducing a manager in a model of random access networks, similar to the model of [14]. Time is divided into slots of equal length, and a user can transmit its packet or wait in each slot. Due to interference, a packet is successfully transmitted only if there is no other packet transmitted in the current slot. If more than one packet is transmitted simultaneously, a collision occurs. The manager can transmit its packets as users do, and it interferes with all the users.

We model the random access scenario as an intervention game. We assume that each user, including the manager, transmits its packets with a fixed probability over time. The action of user i, a_i, is thus its transmission probability, and we have $A_i = [0, 1]$ for all $i \in \mathcal{N}_0$. The average data rate for user $i \in \mathcal{N}$

[6] We define $(\alpha, \beta) = \emptyset$ if $\alpha \geq \beta$.

when the users transmit according to the probabilities $(a_0, a) \in A_0 \times A$ is given by

$$r_i(a_0, a) = \gamma_i a_i \prod_{j \in \mathcal{N}_0 \setminus \{i\}} (1 - a_j),$$

where $\gamma_i > 0$ is the fixed peak data rate for user i. The benefit that a regular user i obtains from its average data rate is represented by a utility function $U_i : \mathbb{R}_+ \to \mathbb{R}$, which is assumed to be strictly increasing. Hence, the utility function of regular user i in the intervention game is given by $u_i(a_0, a) = U_i(r_i(a_0, a))$. We assume that the manager is benevolent with the utility function $u_0(a_0, a) = \sum_{i \in \mathcal{N}} u_i(a_0, a)$ and that its monitoring technology is perfect.

In the above intervention game, the minimal and maximal intervention actions are given by $\underline{a}_0 = 0$ and $\overline{a}_0 = 1$, respectively. Since $r_i(0, a) \geq 0$ and $r_i(1, a) = 0$ for all $a \in A$, for all $i \in \mathcal{N}$, we have $\mathcal{E} = A$ by Proposition 1. In other words, any action profile $a \in A$ can be supported by an intervention mechanism. Because the maximal intervention action yields zero rate to all the users regardless of the action profile, the maximum punishment is strong enough to prevent deviations from any target action profile.[7] Since $A = [0, 1]^N$ is compact and u_0 is continuous, a solution to $\max_{a \in A} u_0(0, a)$ exists. Then by Proposition 2, (f_{a^*}, a^*) is an intervention equilibrium if and only if a^* maximizes $u_0(0, a)$ on A. Also, we can apply Proposition 4 to show that any $a^* \in (0, 1)^N$ is supported by f_{a^*,c^*} with $c_i^* = 1/a_i^*$ for all $i \in \mathcal{N}$. Suppose for the moment that the utility of each user is given by its average data rate, i.e., $u_i(a_0, a) = r_i(a_0, a)$ for all $i \in \mathcal{N}$. Then for each $i \in \mathcal{N}$, we obtain $c_i^* = 1/a_i^* > 0$ by (6), and we can verify that the conditions (8)–(10) are satisfied. Thus, by Proposition 4, we can conclude that f_{a^*,c^*} supports a^*. Note that the concept of an intervention mechanism supporting an action profile is based on Nash equilibrium, which uses only the ordinal properties of the utility functions. Therefore, f_{a^*,c^*} still supports a^* even when the utility function of user i is given by $u_i(a_0, a) = U_i(r_i(a_0, a))$ for any strictly increasing function U_i, for all $i \in \mathcal{N}$.

The above argument points out an informational advantage of direct intervention over indirect intervention. To highlight the informational advantage of direct intervention, suppose that the objective of the manager is to implement a target action profile $a^* \in (0, 1)^N$, determined independently of $(U_1, \ldots, U_N)$, while taking the minimal intervention action when the users choose a^*. Then the results in the previous paragraph imply that the direct intervention mechanisms f_{a^*} and f_{a^*,c^*} with $c_i^* = 1/a_i^*$ for all $i \in \mathcal{N}$ support the action profile a^* for *any* $(U_1, \ldots, U_N)$. Since direct intervention affects utility through its

[7] In fact, a maximum punishment intervention mechanism can prevent not only unilateral deviations from its target action profile but also joint deviations.

impact on the rates, its effectiveness does not depend on the shapes of the utility functions. Thus, the manager with direct intervention does not need to know the utility functions of the users, $(U_1, \ldots, U_N)$, in order to design an intervention mechanism that supports a target action profile.[8] This property can be considered as the robustness of direct intervention with respect to the utilities of the users.

To draw a contrast, consider an alternative intervention scenario where the manager intervenes through pricing. In such a scenario, the action profile of the regular users determines their average data rates, i.e., $r_i(a) = \gamma_i a_i \prod_{j \in \mathcal{N} \setminus \{i\}} (1 - a_j)$ for all $i \in \mathcal{N}$, while a pricing mechanism specifies the payments that the regular users make depending on their action profile, i.e., $f(a) = (f_1(a), \ldots, f_N(a))$, where $f_i(a)$ is the payment of user i. Thus, the utility function of user i is given by

$$u_i(f(a), a) = U_i(r_i(a)) - f_i(a).$$

As can be seen from the above expression, pricing affects utility by taking away utility units from the users, and thus the shapes and scales of the utility functions matter to the manager when designing optimal pricing mechanisms. For example, consider a pricing mechanism that charges each user an amount proportional to its average data rate, i.e., $f_i(a) = p_i r_i(a)$ for all $i \in \mathcal{N}$, where p_i is the price of unit data rate for user i. Then the pricing mechanism supports an action profile a^* when the manager sets $p_i = U_i'(r_i(a^*))$ for all $i \in \mathcal{N}$, assuming that U_i is differentiable and concave for all $i \in \mathcal{N}$. This example illustrates that, in contrast to direct intervention, the manager needs to know the utility functions of the users, $(U_1, \ldots, U_N)$, in order to design a pricing mechanism that supports a target action profile.

VI. Conclusion

This paper presents the intervention framework that is aimed to optimize the objective of the manager, or the network performance, taking the incentives of selfish users into account. We have highlighted the generality and advantages of our framework. In particular, we have pointed out the advantages of intervention over pricing in terms of implementation and informational requirement. To facilitate analysis in the intervention framework, we have developed a new class of games called intervention games and its solution concept called intervention equilibrium. Our analytic results in this paper are limited to the special case of a benevolent manager with perfect monitoring. It is our plan for future research to investigate intervention mechanisms in the case of imperfect monitoring, analyzing how errors in signals affect optimal intervention mechanisms and the network performance.

[8] Of course, if the target action profile depends on $(U_1, \ldots, U_N)$, the manager needs to know $(U_1, \ldots, U_N)$ to determine it.

APPENDIX

PROOF OF PROPOSITION 4

Proof: To prove that f_{a^*,c^*} supports a^* is equivalent to show

$$a_i^* \in \arg\max_{a_i \in A_i} u_i(f_{a^*,c^*}(a_i, a_{-i}^*), a_i, a_{-i}^*) \tag{12}$$

for all $i \in \mathcal{N}$. Note that $f_{a^*,c^*}(a_i, a_{-i}^*) = [c_i^*(a_i - a_i^*) + \underline{a}_0]_{\underline{a}_0}^{\overline{a}_0}$. We consider three cases depending on the sign of c_i^*.

Case 1: $c_i^* = 0$.

By (6), we have $\partial u_i(\underline{a}_0, a^*)/\partial a_i = 0$. Also, we have $f_{a^*,c^*}(a_i, a_{-i}^*) = \underline{a}_0$ for all $a_i \in A_i$. Thus, the objective function in (12) reduces to $u_i(\underline{a}_0, a_i, a_{-i}^*)$. The condition (7) implies that $u_i(\underline{a}_0, a_i, a_{-i}^*)$ is a concave function with respect to a_i on A_i. Also, the first-order optimality condition is satisfied at $a_i = a_i^*$ since $\partial u_i(\underline{a}_0, a^*)/\partial a_i = 0$. Therefore, a_i^* maximizes $u_i(\underline{a}_0, a_i, a_{-i}^*)$ on A_i.

Case 2: $c_i^* > 0$.

Since $\partial u_i(\underline{a}_0, a^*)/\partial a_0 < 0$, we have $\partial u_i(\underline{a}_0, a^*)/\partial a_i > 0$ by (6). First, consider $a_i \in [\underline{a}_i, a_i^*]$. In this region, $f_{a^*,c^*}(a_i, a_{-i}^*) = \underline{a}_0$, and thus the objective function can be written as $u_i(\underline{a}_0, a_i, a_{-i}^*)$. Since the condition (8) implies that $u_i(\underline{a}_0, a_i, a_{-i}^*)$ is concave with respect to a_i on $[\underline{a}_i, a_i^*]$, $u_i(\underline{a}_0, a_i, a_{-i}^*)$ is strictly increasing in a_i on $[\underline{a}_i, a_i^*]$.

Second, consider $a_i \in [a_i^*, \min\{\overline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*\}]$. In this region, $f_{a^*,c^*}(a_i, a_{-i}^*) = c_i^*(a_i - a_i^*) + \underline{a}_0$, and thus the objective function can be written as $u_i(c_i^*(a_i - a_i^*) + \underline{a}_0, a_i, a_{-i}^*)$. The first derivative of $u_i(c_i^*(a_i - a_i^*) + \underline{a}_0, a_i, a_{-i}^*)$ with respect to a_i is given by

$$\left(c_i^* \frac{\partial u_i}{\partial a_0} + \frac{\partial u_i}{\partial a_i} \right) \bigg|_{(a_0, a_i, a_{-i}) = (c_i^*(a_i - a_i^*) + \underline{a}_0, a_i, a_{-i}^*)},$$

while the second derivative is given by the left-hand side of (9). The first derivative is zero at $a_i = a_i^*$ by (6), while the second derivative is non-positive by (9). Hence, the first derivative is non-positive on $(a_i^*, \min\{\overline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*\})$, and thus $u_i(c_i^*(a_i - a_i^*) + \underline{a}_0, a_i, a_{-i}^*)$ is non-increasing in a_i on $[a_i^*, \min\{\overline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*\}]$.

Lastly, consider $a_i \in [a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*, \overline{a}_i]$. In this region, $f_{a^*,c^*}(a_i, a_{-i}^*) = \overline{a}_0$, and thus the objective function can be written as $u_i(\overline{a}_0, a_i, a_{-i}^*)$. Since the first derivative of $u_i(\overline{a}_0, a_i, a_{-i}^*)$ with respect to a_i is non-positive on $(a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*, \overline{a}_i)$ by (10), $u_i(\overline{a}_0, a_i, a_{-i}^*)$ is non-increasing in a_i on $[a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*, \overline{a}_i]$.

Case 3: $c_i^* < 0$.

In this case, $\partial u_i(\underline{a}_0, a^*)/\partial a_i < 0$ and

$$f_{a^*,c^*}(a_i, a_{-i}^*) = \begin{cases} \overline{a}_0 & \text{for } a_i \in [\underline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*], \\ c_i^*(a_i - a_i^*) + \underline{a}_0 & \text{for } a_i \in [\max\{\overline{a}_i, a_i^* + (\overline{a}_0 - \underline{a}_0)/c_i^*\}, a_i^*], \\ \underline{a}_0 & \text{for } a_i \in [a_i^*, \overline{a}_i]. \end{cases}$$

Following an analogous argument as in Case 2, we can show that the objective function is non-decreasing in a_i on $[\underline{a}_i, a_i^*]$ and strictly decreasing on $[a_i^*, \overline{a}_i]$, implying that $a_i = a_i^*$ maximizes the objective function on A_i.

Note that if the inequalities in (7), (9), and (11) are strict, we have $a_i = a_i^*$ as a unique maximizer for all $i \in \mathcal{N}$. ■

REFERENCES

[1] R. Garg, A. Kamra, and V. Khurana, "A game-theoretic approach towards congestion control in communication networks," *Comput. Commun. Review*, vol. 32, no. 3, pp. 47–61, Jul. 2002.

[2] J. Park and M. van der Schaar, "Stackelberg contention games in multiuser networks," *EURASIP J. Advances Signal Process.*, vol. 2009, Article ID 305978, 15 pages, 2009.

[3] Y. A. Korilis, A. A. Lazar, and A. Orda, "Achieving network optima using Stackelberg routing strategies," *IEEE/ACM Trans. Netw.*, vol. 5, no. 1, pp. 161–173, Feb. 1997.

[4] R. T. Ma, V. Misra, and D. Rubenstein, "An analysis of generalized slotted-Aloha protocols," *IEEE/ACM Trans. Netw.*, vol. 17, no. 3, pp. 936–949, Jun. 2009.

[5] Y. Su and M. van der Schaar, "A new perspective on multi-user power control games in interference channels," *IEEE Trans. Wireless Commun.*, vol. 8, no. 6, pp. 2910–2919, Jun. 2009.

[6] J. K. MacKie-Mason and H. R. Varian, "Pricing congestible network resources," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 7, pp. 1141–1149, Sep. 1995.

[7] T. Başar and R. Srikant, "A Stackelberg network game with a large number of followers," *J. Optimization Theory Applicat.*, vol. 115, no. 3, pp. 479–490, Dec. 2002.

[8] M. Bloem, T. Alpcan, and T. Başar, "A Stackelberg game for power control and channel allocation in cognitive radio networks," in *Proc. ValueTools*, 2007.

[9] S. H. Low and D. E. Lapsley, "Optimization flow control—I: Basic algorithm and convergence," *IEEE/ACM Trans. Netw.*, vol. 7, no. 6, pp. 861–874, Dec. 1999.

[10] G. Mailath and L. Samuelson, *Repeated Games and Reputations: Long-run Relationships*. Oxford, U.K.: Oxford Univ. Press, 2006.

[11] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford, U.K.: Oxford Univ. Press, 1995.

[12] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, Eds., *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.

[13] R. Myerson, *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard Univ. Press, 1991.

[14] A. H. Mohsenian-Rad, J. Huang, M. Chiang, and V. W. S. Wong, "Utility-optimal random access without message passing," *IEEE Trans. Wireless Commun.*, vol. 8, no. 3, pp. 1073–1079, Mar. 2009.